

Community-Wide Evaluation of Computational Function Prediction

Iddo Friedberg and Predrag Radivojac

Abstract

A biological experiment is the most reliable way of assigning function to a protein. However, in the era of high-throughput sequencing, scientists are unable to carry out experiments to determine the function of every single gene product. Therefore, to gain insights into the activity of these molecules and guide experiments, we must rely on computational means to functionally annotate the majority of sequence data. To understand how well these algorithms perform, we have established a challenge involving a broad scientific community in which we evaluate different annotation methods according to their ability to predict the associations between previously unannotated protein sequences and Gene Ontology terms. Here we discuss the rationale, benefits, and issues associated with evaluating computational methods in an ongoing community-wide challenge.

Key words Function prediction, Algorithms, Evaluation, Machine learning

1 Introduction

Molecular biology has become a high volume information science. This rapid transformation has taken place over the past two decades and has been chiefly enabled by two technological advances: (1) affordable and accessible high-throughput sequencing platforms, sequence diagnostic platforms, and proteomic platforms and (2) affordable and accessible computing platforms for managing and analyzing these data. It is estimated that sequence data accumulates at the rate of 100 exabases per day (1 exabase = 10^{18} bases) [35]. However, the available sequence data are of limited use without understanding their biological implications. Therefore, the development of computational methods that provide clues about functional roles of biological macromolecules is of primary importance.

Many function prediction methods have been developed over the past two decades [12, 31]. Some are based on sequence alignments to proteins for which the function has been experimentally established [4, 11, 24], yet others exploit other types of data such as

protein structure [26, 27], protein and gene expression data [17], macromolecular interactions [21, 25], scientific literature [3], or a combination of several data types [9, 34, 36]. Typically, each new method is trained and evaluated on different data. Therefore, establishing best practices in method development and evaluating the accuracy of these methods in a standardized and unbiased setting is important. To help choose an appropriate method for a particular task, scientists often form community challenges for evaluating methods [7]. The scope of these challenges extends beyond testing methods: they have been successful in invigorating their respective fields of research by building communities and producing new ideas and collaborations (e.g., [20]).

In this chapter we discuss a community-wide effort whose goal is to help understand the state of affairs in computational protein function prediction and drive the field forward. We are holding a series of challenges which we named the Critical Assessment of Functional Annotation, or CAFA. CAFA was first held in 2010–2011 (CAFA1) and included 23 groups from 14 countries who entered 54 computational function prediction methods that were assessed for their accuracy. To the best of our knowledge, this was the first large-scale effort to provide insights into the strengths and weaknesses of protein function prediction software in the bioinformatics community. CAFA2 was held in 2013–2014, and more than doubled the number of groups (56) and participating methods (126). Although several repetitions of the CAFA challenge would likely give accurate trajectory of the field, there are valuable lessons already learned from the two CAFA efforts.

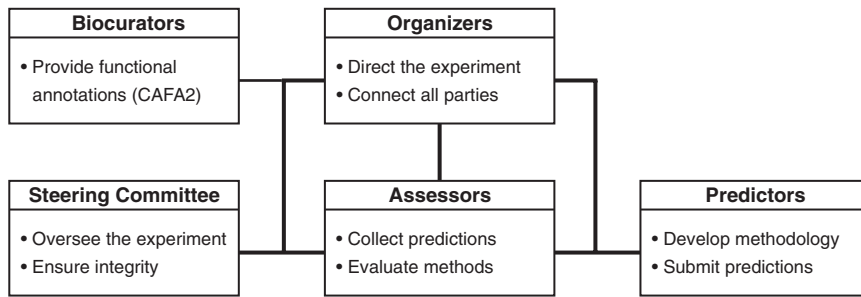
For further reading on CAFA1, the results were reported in full in [30]. As of this time, the results of CAFA2 are still unpublished and will be reported in the near future. The preprint of the paper is available on arXiv [19].

2 Organization of the CAFA Challenge

We begin our explanation of CAFA by describing the participants. The CAFA challenge generally involves the following groups: the organizers, the assessors, the biocurators, the steering committee, and the predictors (Fig. 1a).

The main role of the organizers is to run CAFA smoothly and efficiently. They advertise the challenge to recruit predictors, coordinate activities with the assessors, report to the steering committee, establish the set of challenges and types of evaluation, and run the CAFA web site and social networks. The organizers also compile CAFA data and coordinate the publication process. The assessors develop assessment rules, write and maintain assessment software, collect the submitted prediction data, assess the data, and present the evaluations to the community. The assessors work together with the organizers and the steering committee on standardizing

a CAFA organization



b Experiment timeline




Fig. 1 The organizational structure of the CAFA experiment. **(a)** Five groups of participants in the experiment together with their main roles. Organizers, assessors, and biocurators cannot participate as predictors. **(b)** Timeline of the experiment

submission formats and developing assessment rules. The biocurators joined the experiment during CAFA2: they provide additional functional annotations that may be particularly interesting for the challenge. The steering committee members are in regular contact with the organizers and assessors. They provide advice and guidance that ensures the quality and integrity of the experiment. Finally, the largest group, the predictors, consists of research groups who develop methods for protein function prediction and submit their predictions for evaluation. The organizers, assessors, and biocurators are not allowed to officially evaluate their own methods in CAFA.

CAFA is run as a timed challenge (Fig. 1b). At time t_0 , a large number of experimentally unannotated proteins are made public by the organizers and the predictors are given several months, until time t_1 , to upload their predictions to the CAFA server. At time t_1 the experiment enters a waiting period of at least several months, during which the experimental annotations are allowed to accumulate in databases such as Swiss-Prot [2] and UniProt-GOA [16]. These newly accumulated annotations are collected at time t_2 and are expected to provide experimental annotations for a subset of original proteins. The performance of participating methods is then analyzed between time points t_2 and t_3 and presented to the community at time t_3 . It is important to mention that unlike some machine learning challenges, CAFA organizers do not provide training data that is required to be used. CAFA, thus, evaluates a combination of biological knowledge, the ability to collect and curate training data, and the ability to develop advanced computational methodology.

We have previously described some of the principles that guide us in organizing CAFA [13]. It is important to mention that CAFA is associated with the Automated Function Prediction Special Interest Group (Function-SIG) that is regularly held at the Intelligent Systems for Molecular Biology (ISMB) conference [37]. These meetings provide a forum for exchanging ideas and communicating research among the participants. Function-SIG also serves as the venue at which CAFA results are initially presented and where the feedback from the community is sought.

3 The Gene Ontology Provides the Functional Repertoire for CAFA

Computational function prediction methods have been reviewed extensively [12, 31] and are also discussed in Chapter 5 [8]. Briefly, a function prediction method can be described as a classifier: an algorithm that is tasked with correctly assigning biological function to a given protein. This task, however, is arbitrarily difficult unless the function comes from a finite, preferably small, set of functional terms. Thus, given an unannotated protein sequence and a set of available functional terms, a predictor is tasked with associating terms to a protein, giving a score (ideally, a probability) to each association.

The Gene Ontology (GO) [1] is a natural choice when looking for a standardized, controlled vocabulary for functional annotation. GO's high adoption rate in the protein annotation community helped ensure CAFA's attractiveness, as many groups were already developing function prediction methods based on GO, or could migrate their methods to GO as the ontology of choice. A second consideration is GO's ongoing maintenance: GO is continuously maintained by the Gene Ontology Consortium, edited and expanded based on ongoing discoveries related to the function of biological macromolecules.

One useful characteristic of the basic GO is that its directed acyclic graph structure can be used to quantify the information provided by the annotation; for details on the GO structure see Chaps. 1 and 3 [14, 15]. Intuitively, this can be explained as follows: the annotation term "Nucleic acid binding" is less specific than "DNA binding" and, therefore, is less informative (or has a lower *information content*). (A more precise definition of information content and its use in GO can be found in [23, 32].) The following question arises: if we know that the protein is annotated with the term "Nucleic acid binding," how can we quantify the additional information provided by the term "DNA binding" or incorrect information provided by the term "RNA binding"? The hierarchical nature of GO is therefore important in determining proper metrics for annotation accuracy. The way this is done will be discussed in Sect. 4.2.

When annotating a protein with one or more GO terms, the association of each GO term with the protein should be described using an Evidence Code (EC), indicating how the annotation is

supported. For example, the Experimental Evidence code (EXP) is used in an annotation to indicate that an experimental assay has been located in the literature, whose results indicate a gene product's function. Other experimental evidence codes include Inferred by Expression Pattern (IEP), Inferred from Genetic Interaction (IGI), and Inferred from Direct Assay (IDA), among others. Computational evidence codes include lines of evidence that were generated by computational analysis, such as orthology (ISO), genomic context (IGC), or identification of key residues (IKR). Evidence codes are not intended to be a measure of trust in the annotation, but rather a measure of provenance for the annotation itself. However, annotations with experimental evidence are regarded as more reliable than computational ones, having a provenance stemming from experimental verification. In CAFA, we treat proteins annotated with experimental evidence codes as a "gold standard" for the purpose of assessing predictions, as explained in the next section. The computational evidence codes are treated as predictions.

From the point of view of a computational challenge, it is important to emphasize that the hierarchical nature of the GO graph leads to the property of consistency or True Path Rule in functional annotation. Consistency means that when annotating a protein with a given GO term, it is automatically annotated with all the ancestors of that term. For example, a valid prediction cannot include "DNA binding" but exclude "Nucleic acid binding" from the ontology because DNA binding implies nucleic acid binding. We say that a prediction is not consistent if it includes a child term, but excludes its parent. In fact, the UniProt resource and other databases do not even list these parent terms from a protein's experimental annotation. If a protein is annotated with several terms, a valid complete annotation will automatically include all parent terms of the given terms, propagated to the root(s) of the ontology. The result is that a protein's annotation can be seen as a consistent sub-graph of GO. Since any computational method effectively chooses one of a vast number of possible consistent sub-graphs as its prediction, the sheer size of the functional repertoire suggests that function prediction is non-trivial.

4 Comparing the Performance of Prediction Methods

In the CAFA challenge, we ask the participants to associate a large number of proteins with GO terms and provide a probability score for each such association. Having associated a set of GO sub-graphs with a given confidence, the next step is to assess how accurate these predictions are. This involves: (1) establishing standards of truth and (2) establishing a set of assessment metrics.

4.1 *Establishing Standards of Truth*

The main challenge to establishing a standard-of-truth set for testing function prediction methods is to find a large set of correctly annotated proteins whose functions were, until recently, unknown.

An obvious choice would be to ask experimental scientists to provide these data from their labs. However, scientists prefer to keep the time between discovery and publication as brief as possible, which means that there is only a small window in which new experimental annotations are not widely known and can be used for assessment. Furthermore, each experimental group has its own “data sequestration window” making it hard to establish a common time for all data providers to sequester their data. Finally, to establish a good statistical baseline for assessing prediction method performance, a large number of prediction targets are needed, which is problematic since most laboratories research one or only a few proteins each. High-throughput experiments, on the other hand, provide a large number of annotations, but those tend to concentrate only on few functions, and generally provide annotations that have a lower information content [32].

Given these constraints, we decided that CAFA would not initially rely on direct communication between the CAFA organizers and experimental scientists to provide new functional data. Instead, CAFA relies primarily on established biocuration activities around the world: we use annotation databases to conduct CAFA as a time-based challenge. To do so, we exploit the following dynamics that occurs in annotation databases: protein annotation databases grow over time. Many proteins that at a given time t_1 do not have experimentally verified annotation, but later, some of proteins may gain experimental annotations, as biocurators add these data into the databases. This subset of proteins that were not experimentally annotated at t_1 , but gained experimental annotations at t_2 , are the ones that we use as a test set during assessment (Fig. 1b). In CAFA1 we reviewed the growth of Swiss-Prot over time and chose 50,000 *target proteins* that had no experimental annotation in the Molecular Function or Biological Process ontologies of GO. At t_2 , out of those 50,000 targets we identified 866 *benchmark proteins*, i.e., targets that gained experimental annotation in the Molecular Function and/or Biological Process ontologies. While a benchmark set of 866 proteins constitutes only 1.7% of the number of original targets, it is a large enough set for assessing performance of prediction methods. To conclude, exploiting the history of the Swiss-Prot database enabled its use as the source for standard-of-truth data for CAFA. In CAFA2, we have also considered experimental annotations from UniProt-GOA [16] and established 3681 benchmark proteins out of 100,000 targets (3.7%).

One criticism of a time-based challenge is that when assessing predictions, we still may not have a full knowledge of a protein’s function. A protein may have gained experimental validation for function f_1 , but it may also have another function, say f_2 , associated with it, which has not been experimentally validated by the time t_2 . A method predicting f_2 may be judged to have made a false-positive prediction, even though it is correct (only we do not know it yet). This problem, known as the “incomplete knowledge problem” or

the “open world problem” [10] is discussed in detail in Chapter 8 [33]. Although the incomplete knowledge problem may impact the accuracy of time-based evaluations, its actual impact in CAFA has not been substantial. There are several reasons for this and are also discussed in, including the robustness of the evaluation metrics used in CAFA, and that the newly added terms may be unexpected and more difficult to predict. The influence of incomplete data and conditions under which it can affect a time-based challenge were investigated and discussed in [18]. Another criticism of CAFA is that the experimental functional annotations are not unbiased because some terms have a much higher frequency than others due to artificial considerations. There are two chief reasons for this bias: first, high-throughput assays typically assign shallow terms to proteins, but being high throughput means they can dominate the experimentally verified annotations in the databases. Second, biomedical research is driven by interest in specific areas of human health, resulting in overrepresentation of health-related functions [32]. Unfortunately, CAFA1 and CAFA2 could not guarantee unbiased evaluation. However, we will expand the challenge in CAFA3 to collect genome-wide experimental evidence for several biological terms. Such an assessment will result in unbiased evaluation on those specific terms.

4.2 Assessment Metrics

When assessing the prediction quality of different methods, two questions come to mind. First, what makes a good prediction? Second, how can one score and rank prediction methods? There is no simple answer to either of these questions. As GO comprises three ontologies that deal with different aspects of biological function, different methods should be ranked separately with respect to how well they perform in Molecular Function, Biological Process, or the Cellular Component ontologies. Some methods are trained to predict only for a subset of any given GO graph. For example, they may only provide predictions of DNA-binding proteins or of mitochondrial-targeted proteins. Furthermore, some methods are trained only on a single species or a subset of species (say, eukaryotes), or using specific types of data such as protein structure, and it does not make sense to test them on benchmark sets for which they were not trained. To address this issue, CAFA scored methods not only in general performance, but also on specific subsets of proteins taken from humans and model organisms, including *Mus musculus*, *Rattus norvegicus*, *Arabidopsis thaliana*, *Drosophila melanogaster*, *Caenorhabditis elegans*, *Saccharomyces cerevisiae*, *Dictyostelium discoideum*, and *Escherichia coli*. In CAFA2, we extended this evaluation to also assess the methods only on benchmark proteins on which they made predictions; i.e., the methods were not penalized for omitting any benchmark protein.

One way to view function prediction is as an information retrieval problem, where the most relevant functional terms should be correctly retrieved from GO and properly assigned to the amino acid sequence at hand. Since each term in the ontology implies

some or all of its ancestors,¹ a function prediction program's task is to assign the best consistent sub-graph of the ontology to each new protein and output a prediction score for this sub-graph and/or each predicted term. An intuitive scoring mechanism for this type of problem is to treat each term independently and provide the precision–recall curve. We chose this evaluation as our main evaluation in CAFA1 and CAFA2.

Let us provide more detail. Consider a single protein on which evaluation is carried out, but keep in mind that CAFA eventually averages all metrics over the set of benchmark proteins. Let now T be a set of experimentally determined nodes and P a non-empty set of predicted nodes in the ontology for the given protein. Precision (pr) and recall (rc) are defined as

$$\text{pr}(P, T) = \frac{|P \cap T|}{|P|}; \quad \text{rc}(P, T) = \frac{|P \cap T|}{|T|},$$

where $|P|$ is the number of predicted terms, $|T|$ is the number of experimentally determined terms, and $|P \cap T|$ is the number of terms appearing in both P and T ; see Fig. 2 for an illustrative example of this measure. Usually, however, methods will associate scores with each predicted term and then a set of terms P will be established by defining a score threshold t , i.e., all predicted terms with scores greater than t will constitute the set P . By varying the decision threshold $t \in [0, 1]$, the precision and recall of each method can be plotted as a curve $(\text{pr}(t), \text{rc}(t))$, where one axis is the precision and the other the recall; see Fig. 3 for an illustration of pr–rc curves and [30] for pr–rc curves in CAFA1. To compile the precision–recall information into a single number that would allow easy comparison between methods, we used the maximum harmonic mean of precision and recall anywhere on the curve, or the maximum F_1 -measure which we call F_{\max}

$$F_{\max} = \max_t \left\{ 2 \times \frac{\text{pr}(t) \times \text{rc}(t)}{\text{pr}(t) + \text{rc}(t)} \right\},$$

where we modified $\text{pr}(t)$ and $\text{rc}(t)$ to reflect the dependency on t . It is worth pointing out that the F-measure used in CAFA places equal emphasis on precision and recall as it is unclear which of the two should be weighted more. One alternative to F_1 would be the use of a combined measure that weighs precision over recall, which reflects the preference of many biologists for few answers with a high fraction of correctly predicted terms (high precision) over many answers with a lower fraction of correct predictions (high recall); the rationale for this tradeoff is illustrated in Fig. 3.

¹Some types of edges in Gene Ontology violate the transitivity property (consistency assumption), but they are not frequent.




Fig. 2 CAFA assessment metrics. (a) *Red nodes* are the predicted terms P for a particular decision threshold in a hypothetical ontology and (b) *blue nodes* are the true, experimentally determined terms T . The *circled terms* represent the overlap between the predicted sub-graph and the true sub-graph. There are two nodes (*circled*) in the intersection of P and T , where $|P| = 5$ and $|T| = 3$. This sets the prediction's precision at $2/5 = 0.4$ and recall at $2/3 = 0.667$, with $F_1 = 2 \times 0.4 \times 0.667 / (0.4 + 0.667) = 0.5$. The remaining uncertainty (*ru*) is the information content of the *uncircled blue node* in panel (b), while the *misinformation (mi)* is the total information content of the *uncircled red nodes* in panel (a). An information content of any node v is calculated from a representative database as $-\log\text{Pr}(v|Pa(v))$; i.e., the probability that the node is present in a protein's annotation given that all its parents are also present in its annotation

However, preferring precision over recall in a hierarchical setting can steer methods to focus on shallow (less informative) terms in the ontology and thus be of limited use. At the same time, putting more emphasis on recall may lead to overprediction, a situation in which many or most of the predicted terms are incorrect. For this reason, we decided to equally weight precision and recall. Additional metrics within the precision–recall framework have been considered, though not implemented yet.

Precision and recall are useful because they are easy to interpret: a precision of $1/2$ means that one half of all predicted terms are correct, whereas a recall of $1/3$ means that one third of the experimental terms have been recovered by the predictor. Unfortunately, precision–recall curves and F_1 , while simple and interpretable measures for evaluating ontology-based predictions, are limited because they ignore the hierarchical nature of the ontology and dependencies among terms. They also do not directly capture the information content of the predicted terms. Assessment metrics that take into account the information content of the terms were developed in the past [22, 23, 29], and are also detailed in Chapter 12 [28]. In CAFA2 we used an information-theoretic measure in which each term is assigned a probability that is dependent on the probabilities of its direct parents. These probabilities are calculated from the frequencies of the terms in the database used to generate the CAFA targets. The entire ontology graph, thus, can be seen as a simple Bayesian network [5]. Using this representation, two information-theoretic analogs of precision and recall can be constructed. We refer to these quantities as *misinformation (mi)*, the information content attributed to the nodes in the predicted graph that are incorrect, and

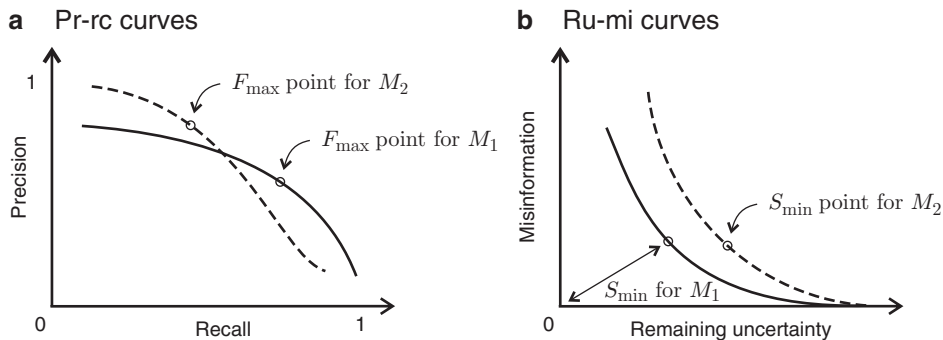


Fig. 3 Precision-recall curves and remaining uncertainty-misinformation curves. This figure illustrates the need for multiple assessment metrics, and understanding the context in which the metrics are used. **(a)** two pr-rc curves corresponding to two prediction methods M_1 and M_2 . The point on each curve that gives F_{\max} is marked as a circle. Although the two methods have a similar performance according to F_{\max} , method M_1 achieves its best performance at high recall values, whereas method M_2 achieves its best performance at high precision values. **(b)** two ru-mi curves corresponding to the same two prediction methods with marked points where the minimum semantic distance is achieved. Although the two methods have similar performance in the pr-rc space, method M_1 outperforms M_2 in ru-mi space. Note, however, that the performance in ru-mi space depends on the frequencies of occurrence of every term in the database. Thus, two methods may score differently in their S_{\min} when the reference database changes over time, or using a different database

remaining uncertainty (ru), the information content of all nodes that belong to the true annotation but not the predicted annotation. More formally, if T is a set of experimentally determined nodes and P a set of predicted nodes in the ontology, then

$$ru(P, T) = - \sum_{v \in T-P} \log \Pr(v | Pa(v)); \quad mi(P, T) = - \sum_{v \in P-T} \log \Pr(v | Pa(v)),$$

where $Pa(v)$ is the set of parent terms of the node v in the ontology (Fig. 2). A single performance measure to rank methods, the minimum semantic distance S_{\min} , is the minimum distance from the origin to the curve $(ru(t), mi(t))_t$. It is defined as

$$S_{\min} = \min_t \left\{ (ru^k(t) + mi^k(t))^{\frac{1}{k}} \right\},$$

where $k \geq 1$. We typically choose $k=2$, in which case S_{\min} is the minimum Euclidean distance between the ru-mi curve and the origin of the coordinate system (Fig. 3b). The ru-mi plots and S_{\min} metrics compare the true and predicted annotation graphs by adding an additional weighting component to high-information nodes. In that manner, predictions with a higher information content will be assigned larger weights. The semantic distance has been a useful measure in CAFA2 as it properly accounts for term dependencies in the ontology. However, this approach also has limitations in that it

relies on an assumed Bayesian network as a generative model of protein function as well as on the available databases of protein functional annotations where term frequencies change over time. While the latter limitation can be remedied by more robust estimation of term frequencies in a large set of organisms, the performance accuracies in this setting are generally less comparable over two different CAFA experiments than in the precision–recall setting.

5 Discussion

Critical assessment challenges have been successfully adopted in a number of fields due to several factors. First, the recognition that improvements to methods are indeed necessary. Second, the ability of the community to mobilize enough of its members to engage in a challenge. Mobilizing a community is not a trivial task, as groups have their own research priorities and only a limited amount of resources to achieve them, which may deter them from undertaking a time-consuming and competitive effort a challenge may pose. At the same time, there are quite a few incentives to join a community challenge. Testing one's method objectively by a third party can establish credibility, help point out flaws, and suggest improvements. Engaging with other groups may lead to collaborations and other opportunities. Finally, the promise of doing well in a challenge can be a strong incentive heralding a group's excellence in their field. Since the assessment metrics are crucial to the performance of the teams, large efforts are made to create multiple metrics and to describe exactly what they measure. Good challenge organizers try to be attentive to the requests of the participants, and to have the rules of the challenge evolve based on the needs of the community. An understanding that a challenge's ultimate goal is to improve methodologies and that it takes several rounds of repeating the challenge to see results.

The first two CAFA challenges helped clarify that protein function prediction is a vibrant field, but also one of the most challenging tasks in computational biology. For example, CAFA provided evidence that the available function prediction algorithms outperform a straightforward use of sequence alignments in function transfer. The performance of methods in the Molecular Function category has consistently been reliable and also showed progress over time (unpublished results from CAFA2). On the other hand, the performance in the Biological Process or Cellular Component ontologies has not yet met expectations. One of the reasons for this may be that the terms in these ontologies are less predictable using amino acid sequence data and instead would rely more on high-quality systems data; e.g., see [6]. The challenge has also helped clarify the problems of evaluation, both in terms of evaluating over consistent sub-graphs in the ontology but also in the presence of incomplete and biased molecular data. Finally, although

it is still early, some best practices in the field are beginning to emerge. Exploiting multiple types of data is typically advantageous, although we have observed that both machine learning expertise and good biological insights tend to result in strong performance. Overall, while the methods in the Molecular Function ontology seem to be maturing, in part because of the strong signal in sequence data, the methods in the Biological Process and Cellular Component ontologies still appear to be in the early stages of development. With the help of better data over time, we expect significant improvements in these categories in the future CAFA experiments.

Overall, CAFA generated a strong positive response to the call for both challenge rounds, with the number of participants substantially growing between CAFA1 (102 participants) and CAFA2 (147). This indicates that there exists significant interest in developing computational protein function prediction methods, in understanding how well they perform, and in improving their performance. In CAFA2 we preserved the experiment rules, ontologies, and metrics we used in CAFA1, but also added new ones to better capture the capabilities of different methods. The CAFA3 experiment will further improve evaluation by facilitating unbiased evaluation for several select functional terms.

More rounds of CAFA are needed to know if computational methods will improve as a direct result of this challenge. But given the community's growth and growing interest, we believe that CAFA is a welcome addition to the community of protein function annotators.

Acknowledgements

We thank Kymberleigh Pagel and Naihui Zhou for helpful discussions. This work was partially supported by NSF grants DBI-1458359 and DBI-1458477. Open Access charges were funded by the University College London Library, the Swiss Institute of Bioinformatics, the Agassiz Foundation, and the Foundation for the University of Lausanne.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work's Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work's Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.

References

- Ashburner M, Ball CA, Blake JA, Botstein D, Butler H, Cherry JM, Davis AP, Dolinski K, Dwight SS, Eppig JT, Harris MA, Hill DP, Issel-Tarver L, Kasarskis A, Lewis S, Matese JC, Richardson JE, Ringwald M, Rubin GM, Sherlock G (2000) Gene ontology: tool for the unification of biology. *Nat Genet* 25(1):25–29.
- Bairoch A, Apweiler R, Wu CH, Barker WC, Boeckmann B, Ferro S, Gasteiger E, Huang H, Lopez R, Magrane M, Martin MJ, Natale DA, O'Donovan C, Redaschi N, Yeh LS (2005) The Universal Protein Resource (UniProt). *Nucleic Acids Res* 33(Database issue):D154–D159
- Camon EB, Barrell DG, Dimmer EC, Lee V, Magrane M, Maslen J, Binns D, Apweiler R (2005) An evaluation of GO annotation retrieval for BioCreAtIvE and GOA. *BMC Bioinformatics* 6(Suppl 1):S17
- Clark WT, Radivojac P (2011) Analysis of protein function and its prediction from amino acid sequence. *Proteins* 79(7):2086–2096
- Clark WT, Radivojac P (2013) Information-theoretic evaluation of predicted ontological annotations. *Bioinformatics* 29(13):i53–i61.
- Costanzo M, Baryshnikova A, Bellay J, Kim Y, Spear ED, Sevier CS, Ding H, Koh JL, Toufighi K, Mostafavi S, Prinz J, St Onge RP, VanderSluis B, Makhnevych T, Vizeacoumar FJ, Alizadeh S, Bahr S, Brost RL, Chen Y, Cokol M, Deshpande R, Li Z, Lin ZY, Liang W, Marback M, Paw J, San Luis BJ, Shuteriqi E, Tong AH, van Dyk N, Wallace IM, Whitney JA, Weirauch MT, Zhong G, Zhu H, Houry WA, Brudno M, Ragibizadeh S, Papp B, Pal C, Roth FP, Giaever G, Nislow C, Troyanskaya OG, Bussey H, Bader GD, Gingras AC, Morris QD, Kim PM, Kaiser CA, Myers CL, Andrews BJ, Boone C (2010) The genetic landscape of a cell. *Science* 327(5964):425–431
- Costello JC, Stolovitzky G (2013) Seeking the wisdom of crowds through challenge-based competitions in biomedical research. *Clin Pharmacol Ther* 93(5):396–398
- Cozzetto D, Jones DT (2016) Computational methods for annotation transfers from sequence. In: Dessimoz C, Škunca N (eds) *The gene ontology handbook. Methods in molecular biology*, vol 1446. Humana Press. Chapter 5
- Cozzetto D, Buchan DWA, Bryson K, Jones DT (2013) Protein function prediction by massive integration of evolutionary analyses and multiple data sources. *BMC Bioinformatics* 14(Suppl 3):S1+.
- Dessimoz C, Škunca N, Thomas PD (2013) CAFA and the open world of protein function predictions. *Trends Genet* 29(11):609–610
- Engelhardt BE, Jordan MI, Muratore KE, Brenner SE (2005) Protein molecular function prediction by Bayesian phylogenomics. *PLoS Comput Biol* 1(5):e45
- Friedberg I (2006) Automated protein function prediction—the genomic challenge. *Brief Bioinform* 7(3):225–242.
- Friedberg I, Wass MN, Mooney SD, Radivojac P (2015) Ten simple rules for a community computational challenge. *PLoS Comput Biol* 11(4):e1004150 (2015)
- Gaudet P, Škunca N, Hu JC, Dessimoz C (2016) Primer on the gene ontology. In: Dessimoz C, Škunca N (eds) *The gene ontology handbook. Methods in molecular biology*, vol 1446. Humana Press. Chapter 3
- Hastings J (2016) Primer on ontologies. In: Dessimoz C, Škunca N (eds) *The gene ontology handbook. Methods in molecular biology*, vol 1446. Humana Press. Chapter 1
- Huntley RP, Sawford T, Mutowo-Meullenet P, Shypitsyna A, Bonilla C, Martin MJ, O'Donovan C (2015) The GOA database: gene ontology annotation updates for 2015. *Nucleic Acids Res* 43(Database issue):D1057–D1063
- Huttenhower C, Hibbs M, Myers C, Troyanskaya OG (2006) A scalable method for integration and functional analysis of multiple microarray datasets. *Bioinformatics* 22(23):2890–2897
- Jiang Y, Clark WT, Friedberg I, Radivojac P (2014) The impact of incomplete knowledge on the evaluation of protein function prediction: a structured-output learning perspective. *Bioinformatics* (Oxford, England) 30(17):i609–i616.
- Jiang Y, Oron TR, Clark WT, Bankapur AR, D'Andrea D, Lepore R, Funk CS, Kahanda I, Verspoor KM, Ben-Hur A, Koo E, Penfold-Brown D, Shasha D, Youngs N, Bonneau R, Lin A, Sahraeian SME, Martelli PL, Profitti G, Casadio R, Cao R, Zhong Z, Cheng J, Altenhoff A, Škunca N, Dessimoz C, Dogan T, Hakala K, Kaewphan S, Mehryary F, Salakoski T, Ginter F, Fang H, Smithers B, Oates M, Gough J, Toronen P, Koskinen P, Holm L, Chen CT, Hsu WL, Bryson K, Cozzetto D, Minneci F, Jones DT, Chapman S, Dukka BKC, Khan IK, Kihara D, Ofer D, Rappoport N, Stern A, Cibrian-Uhalte E, Denny P, Foulger RE, Hieta R, Legge D, Lovering RC, Magrane M, Melidoni AN, Mutowo-Meullenet P, Pichler K, Shypitsyna A, Li B, Zakeri P, ElShal S, Tranchevent LC, Das S, Dawson NL, Lee D, Lees JG, Sillitoe I, Bhat P, Nepusz T, Romero AE, Sasidharan R, Yang H, Paccanaro A, Gillis J, Sedenio-Cortes AE, Pavlidis P, Feng S, Cejuela JM, Goldberg T, Hamp T, Richter L, Salamov A, Gabaldon T, Marcet-Houben M, Supek F, Gong Q, Ning W, Zhou Y, Tian W, Falda M, Fontana P, Lavezzo E, Toppo S, Ferrari C, Giollo M, Piovesan D,

- Tosatto S, del Pozo A, Fernández JM, Maietta P, Valencia A, Tress ML, Benso A, Di Carlo S, Politano G, Savino A, Ur Rehman H, Re M, Mesiti M, Valentini G, Bargsten JW, van Dijk ADJ, Gemovic B, Glisic S, Perovic V, Veljkovic V, Veljkovic N, Almeida-e Silva DC, Vencio RZN, Sharan M, Vogel J, Kansakar L, Zhang S, Vucetic S, Wang Z, Sternberg MJE, Wass MN, Huntley RP, Martin MJ, O'Donovan C, Robinson PN, Moreau Y, Tramontano A, Babbitt PC, Brenner SE, Linal M, Orengo CA, Rost B, Greene CS, Mooney SD, Friedberg I, Radivojac P (2016) An expanded evaluation of protein function prediction methods shows an improvement in accuracy. <http://arxiv.org/abs/1601.00891>
20. Kryshchafovych A, Fidelis K, Moulton J (2014) CASP10 results compared to those of previous CASP experiments. *Proteins* 82: 164–174.
 21. Letovsky S, Kasif S (2003) Predicting protein function from protein/protein interaction data: a probabilistic approach. *Bioinformatics* 19(Suppl 1):i197–204
 22. Lord PW, Stevens RD, Brass A, Goble CA (2003) Investigating semantic similarity measures across the gene ontology: the relationship between sequence and annotation. *Bioinformatics* 19(10):1275–1283.
 23. Lord PW, Stevens RD, Brass A, Goble CA (2003) Semantic similarity measures as tools for exploring the gene ontology. In: Pacific symposium on biocomputing. Pacific symposium on biocomputing, pp 601–612.
 24. Martin DM, Berriman M, Barton GJ (2004) GOTcha: a new method for prediction of protein function assessed by the annotation of seven genomes. *BMC Bioinformatics* 5:178
 25. Nabieva E, Jim K, Agarwal A, Chazelle B, Singh M (2005) Whole-proteome prediction of protein function via graph-theoretic analysis of interaction maps. *Bioinformatics* 21(Suppl 1):i302–i310
 26. Pal D, Eisenberg D (2005) Inference of protein function from protein structure. *Structure* 13(1):121–130 (2005)
 27. Pazos F, Sternberg MJ (2004) Automated prediction of protein function and detection of functional sites from structure. *Proc Natl Acad Sci USA* 101(41):14754–14759
 28. Pesquita C (2016) Semantic Similarity in the Gene Ontology. In: Dessimoz C, Škunca N (eds) *The gene ontology handbook. Methods in molecular biology*, vol 1446. Humana Press. Chapter 12
 29. Pesquita C, Faria D, Falcão AO, Lord P, Couto FM (2009) Semantic similarity in biomedical ontologies. *PLoS Comput Biol* 5(7): e1000443+.
 30. Radivojac P, Clark WT, Oron TRR, Schnoes AM, Wittkop T, Sokolov A, Graim K, Funk C, Verspoor K, Ben-Hur A, Pandey G, Yunes JM, Talwalkar AS, Repo S, Souza ML, Piovesan D, Casadio R, Wang Z, Cheng J, Fang H, Gough J, Koskinen P, Törönen P, Nokso-Koivisto J, Holm L, Cozzetto D, Buchan DW, Bryson K, Jones DT, Limaye B, Inamdar H, Datta A, Manjari SK, Joshi R, Chitale M, Kihara D, Lisewski AM, Erdin S, Venner E, Lichtarge O, Rentzsch R, Yang H, Romero AE, Bhat P, Paccanaro A, Hamp T, Kaßner R, Seemayer S, Vicedo E, Schaefer C, Achten D, Auer F, Boehm A, Braun T, Hecht M, Heron M, Hönigschmid P, Hopf TA, Kaufmann S, Kiening M, Krompass D, Landerer C, Mahlich Y, Roos M, Björne J, Salakoski T, Wong A, Shatkay H, Gatzmann F, Sommer I, Wass MN, Sternberg MJ, Škunca N, Supek F, Bošnjak M, Panov P, Džeroski S, Šmuc T, Kourmpetis YA, van Dijk AD, ter Braak CJ, Zhou Y, Gong Q, Dong X, Tian W, Falda M, Fontana P, Lavezzo E, Di Camillo B, Toppo S, Lan L, Djuric N, Guo Y, Vucetic S, Bairoch A, Linal M, Babbitt PC, Brenner SE, Orengo C, Rost B, Mooney SD, Friedberg I (2013) A large-scale evaluation of computational protein function prediction. *Nat Methods* 10(3):221–227.
 31. Rentzsch R, Orengo CA (2009) Protein function prediction—the power of multiplicity. *Trends Biotechnol* 27(4):210–219.
 32. Schnoes AM, Ream DC, Thorman AW, Babbitt PC, Friedberg I (2013) Biases in the experimental annotations of protein function and their effect on our understanding of protein function space. *PLoS Comput Biol* 9(5):e1003063+.
 33. Škunca N, Roberts RJ, Steffen M (2016) Evaluating computational gene ontology annotations. In: Dessimoz C, Škunca N (eds) *The gene ontology handbook. Methods in molecular biology*, vol 1446. Humana Press. Chapter 8.
 34. Sokolov A, Ben-Hur A (2010) Hierarchical classification of gene ontology terms using the GOstruct method. *J Bioinform Comput Biol* 8(2):357–376
 35. Stephens ZD, Lee SY, Faghri F, Campbell RH, Zhai C, Efron MJ, Iyer R, Schatz MC, Sinha S, Robinson GE (2015) Big data: astronomical or genomics? *PLoS Biol* 13(7):e1002195+.
 36. Troyanskaya OG, Dolinski K, Owen AB, Altman RB, Botstein D (2003) A Bayesian framework for combining heterogeneous data sources for gene function prediction (in *Saccharomyces cerevisiae*). *Proc Natl Acad Sci USA* 100(14):8348–8353
 37. Wass MN, Mooney SD, Linal M, Radivojac P, Friedberg I (2014) The automated function prediction SIG looks back at 2013 and prepares for 2014. *Bioinformatics (Oxford, England)* 30(14):2091–2092.